Efficient Authentication of Outsourced String Similarity Search

Boxiang Dong, Wendy Wang
Department of Computer Science
Stevens Institute of Technology
Hoboken, NJ
bdong, hwang4@stevens.edu

arXiv:1603.02727v1 [cs.CR] 8 Mar 2016

ABSTRACT

Cloud computing enables the outsourcing of big data analytics, where a third-party server is responsible for data storage and processing. In this paper, we consider the outsourcing model that provides string similarity search as the service. In particular, given a similarity search query, the service provider returns all strings from the outsourced dataset that are similar to the query string. A major security concern of the outsourcing paradigm is to authenticate whether the service provider returns *sound* and *complete* search results. In this paper, we design $AutoS^3$, an authentication mechanism of outsourced string similarity search. The key idea of $AutoS^3$ is that the server returns a *verification object* (VO) to prove the result correctness. First, we design an authenticated string indexing structure named MB-tree for VO construction. Second, we design two lightweight authentication methods named VS^2 and E-VS^2 that can catch the service provider's various cheating behaviors with cheap verification cost. Moreover, we generalize our solution for top-k string similarity search. We perform an extensive set of experiment results on real-world datasets to demonstrate the efficiency of our approach.

1. INTRODUCTION

Big data analytics offers the promise of providing valuable insights. However, many companies, especially the small- and medium-sized organizations lack the computational resources, in-house knowledge and experience of big data analytics. A practical solution to this dilemma is outsourcing, where the data owner outsources the data to a computational powerful third-party service provider (e.g., the cloud) for cost-effective solutions of data storage, processing, and analysis.

In this paper, we consider *string similarity search*, an important data analytics operation that have been used in a broad range of applications, as the outsourced computations. Generally speaking, the data owner outsources a string database D to a third-party service provider (server). The server provides the storage and processing of similarity search queries as services. The search queries ask for the strings in D that are similar to a number of given strings, where the similarity is measured by a specific similarity function and a user-defined threshold.

For all the benefits of outsourcing and cloud computing, though, the outsourcing paradigm deprives the data owner of direct control over her data. This poses numerous security challenges. One of the challenges is that the server may cheat on the similarity search results. For example, the server is incentivized to improve its revenue by computing with less resources (e.g., only search a portion of D) while charging for more. Therefore, it is important to authenticate whether the service provider has performed the search faithfully, and returned the correct results to the client. A naive method is to execute the search queries locally, and compare the results with the outcome from the server. Apparently this method is prohibitively costly. We aim to design efficient methods that enable the client to authenticate that the server returned *sound* and *complete* similar strings. By soundness we mean that the returned strings are indeed similar. By completeness we mean that all similar strings are returned. In this paper, we focus on *edit distance*, a commonly-used string similarity function.

Most existing work (e.g. [3, 15]) solve the authentication problem for spatial queries in the Euclidean space. To our best knowledge, ours is the first to consider the authentication of outsourced string similarity search. Intuitively, the strings can be mapped to the Euclidean space via a similarity-preserving embedding function (e.g. [8, 13]). However, such embedding functions cannot guarantee 100% precision (i.e., the embedded points of some dissimilar strings become similar in the Euclidean space). This disables the direct use of the existing Euclidean distance based authentication approaches on string similarity queries.

In this paper, we design $AutoS^3$, an Authentication mechanism of Outsourced String Similarity Search. The key idea of $AutoS^3$ is that besides returning the similar strings, the server returns a verification object (VO) that can prove the soundness and completeness of returned strings. In particular, we make the following contributions.

First, we design an authentication tree structure named *MB-tree*. MB-tree is constructed by integrating Merkle hash tree [28], a popularly-used authenticated data structure, with B^{ed}-tree [33], a compact index for efficient string similarity search based on edit distance.

Second, we design the basic verification method named VS^2 for the search queries that consist of a single query

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.
Copyright 20XX ACM X-XXXXX-XX-X/XX/XX ...$15.00.

string. VS^2 constructs VO from the MB-tree, requiring to include *false hits* into VO, where false hits refer to the strings that are not returned in the result, but are necessary for the result authentication. We prove that VS^2 is able to catch the server's cheating behaviors such as tampered values, soundness violation, and completeness violation.

A large amounts of false hits can impose a significant burden to the client for verification. Therefore, our third contribution is the design of the *E-*VS^2 algorithm that reduces the VO verification cost at the client side. E-VS^2 applies a similarity-preserving embedding function to map strings to the Euclidean space in the way that similar strings are mapped to close Euclidean points. Then VO is constructed from both the MB-tree and the embedded Euclidean space. Compared with VS^2, E-VS^2 dramatically saves the verification cost by replacing a large amounts of expensive string edit distance calculation with a small number of cheap Euclidean distance computation.

Fourth, we extend to the authentication of: (1) similarity search queries that consists of multiple query strings, and (2) top-k similarity search. We design efficient optimization methods that reduce verification cost for both cases.

Last but not least, we complement the theoretical investigation with a rich set of experiment study on real datasets. The experiment results demonstrate the efficiency of our approaches. It shows that E-VS^2 can save 25% verification cost of the VS^2 approach.

The rest of the paper is organized as follows. Sections 2 and 3 discuss the related work and preliminaries. Section 4 formally defines the problem. Section 5 presents our VS^2 and E-VS^2 approaches for single-string search queries. Section 6 discusses the authentication of multi-string search queries. Section 7 extends to top-k similarity search. The experiment results are shown in Section 8. Section 9 concludes the paper.

2. RELATED WORK

The problem of authentication of outsourced computations caught much attention from the research community in recent years. Based on the type of the outsourced data and the type of queries that are executed on the data, we classify these techniques into the following types: (1) authentication of outsourced SQL query evaluation, (2) authentication of keyword search, and (3) authentication of outsourced spatial query evaluation. None of these work considered string similarity search queries.

Authentication of outsourced SQL query evaluation. The issue of providing authenticity for outsourced database was initially raised in the database-as-a-service (DaS) paradigm [11]. The aim is to assure the correctness of SQL query evaluation over the outsourced databases. The proposed solutions include Merkle hash trees [19, 23], signatures on a chain of paired tuples [25], and authenticated B-tree and R-tree structures for aggregated queries [21]. The key idea of these techniques is that the data owner outsources not only data but also the endorsements of the data being outsourced. These endorsements are signed by the data owner against tampering with by the service provider. For the cleaning results, the service provider returns both the results and a proof, called the *verification object* (VO), which is an auxiliary data structure to store the processing traces such as index traversals. The client uses the VO, together with the answers, to reconstruct the endorsements and thus verify the authenticity of the results. An efficient authentication technique should minimize the size of VO, while requiring lightweight authentication at the client side. In this paper, we follow the same VO-based strategy to design our authentication method.

Authentication of keyword search. Pang et al. [24] targeted at search engines that perform similarity-based document retrieval, and designed a novel authentication mechanism for the search results. The key idea of the authentication is to build the Merkle hash tree (MHT) on the inverted index, and use the MHT for VO construction. Though effective, it has several limitations, e.g., the MHT cannot deal data updates efficiently [10]. To address these limitations, Goodrich et al. [10] designed a new model that considered conjunctive keyword searches as equivalent with a set intersection on the underlying inverted index data structure. They use the authenticated data structure in [27] to verify the correctness of set operations.

Authentication of outsourced spatial query evaluation. A number of work investigated the problem of authentication of spatial query evaluation in the location-based service model. Yang et al. [31, 30] integrated an R-tree with the MHT (which is called Merkle R-tree or MR-tree) for authenticating multi-dimensional range queries. Yiu et al. [32] focused on the moving KNN queries that continuously reports the k nearest neighbors of a moving query point. They designed the Voronoi MR-tree as the authenticated data structure for VO construction and authentication. Hu et al. [15] also utilized neighborhood information derived from the Voronoi diagram of the underlying spatial dataset for authentication. Wu et al. [29] designed a novel authenticated data structure named *Merkle-IR-tree* (MIR-tree) for moving top-k spatial keyword (MkSK) queries. MIR-tree builds a series of digests in each node of the IR-tree [6], in which each entry summarizes the spatial distances and text relevance of the entries in its child node.

3. PRELIMINARIES

String similarity function. String similarity search is a fundamental problem in many research areas, e.g., information retrieval, database joins, and more. In the literature, there are a number of string similarity functions, e.g., Hamming distance, n-grams, and edit distance (See [17] for a good tutorial.) In this paper, we mainly consider *edit distance*, one of the most popular string similarity measurement that has been used in a wide spectrum of applications. Informally, the edit distance of two string values s_1 and s_2, denoted as $DST(s_1, s_2)$, measures the minimum number of insertion, deletion and substitution operations to transform s_1 to s_2. We say two strings s_1 and s_2 are θ*-similar*, denoted as $s_1 \approx s_2$, if $DST(s_1, s_2) \leq \theta$, where θ is a user-specified similarity threshold. Otherwise, we say s_1 and s_2 are θ-*dissimilar* (denoted as $s_1 \not\approx s_2$).

Mapping Strings to Euclidean Space. Given two strings s_1 and s_2, normally the complexity of computing edit distance is $O(|s_1||s_2|)$, where $|s_1|$ and $|s_2|$ are the lengths of s_1 and s_2. One way to reduce the complexity of similarity measurement is to map the strings into a multi-dimensional Euclidean space, such that the similar strings are mapped to close Euclidean points. The main reason of the embedding is that the computation of Euclidean distance is much cheaper than string edit distance. A few string embedding techniques (e.g., [8, 16, 13]) exist in the literature. These al-

gorithms have different properties in terms of their efficiency and distortion rate (See [12] for a good survey). In this paper, we consider an important property named *contractiveness* property of the embedding methods, which requires that for any pair of strings (s_i, s_j) and their embedded Euclidean points (p_i, p_j), $dst(p_i, p_j) \leq DST(s_i, s_j)$, where $dst()$ and $DST()$ are the distance function in the Euclidean space and string space respectively. In this paper, we use $dst()$ and $DST()$ to denote the Euclidean distance and edit distance. In this paper, we use the *SparseMap* method [13] for the string embedding. SparseMap preserves the contractiveness property. We will show how to leverage the contractiveness property to improve the performance of verification in Section 5.2. Note that the embedding methods may introduce false positives, i.e. the embedding points of dissimilar strings may become close in the Euclidean space.

Authenticated data structure. To enable the client to verify the correctness of the query results, the server returns the results along with some supplementary information that permits result verification. Normally the supplementary information takes the format of *verification objects* (VOs). In the literature, VO generation is usually performed by an authenticated data structure (e.g., [20, 30, 26]). One of the popular authenticated data structures is *Merkle tree* [28]. In particular, a Merkle tree is a tree T in which each leaf node N stores the digest of a record r: $h_N = h(r)$, where $h()$ is a one-way, collision-resistant hash function (e.g. SHA-1 [22]). For each non-leaf node N of T, it is assigned the value $h_N = h(h_{C_1}||\ldots||h_{C_k})$, where $C_1, \ldots, C_k$ are the children of N. The root signature sig is generated by signing the digest h_{root} of the root node using the private key of a trusted party (e.g., the data owner). The VO enables the client to re-construct the root signature sig.

B^{ed}-Tree for string similarity search. A number of compact data structures (e.g., [1, 2, 18]) are designed to handle edit distance based similarity measurement. In this paper, we consider B^{ed}-*tree* [33] due to its simplicity and efficiency. B^{ed}-tree is a B^+-tree based index structure that can handle arbitrary edit distance thresholds. The tree is built upon a *string ordering* scheme which is a mapping function φ to map each string to an integer value. To simplify notation, we say that $s \in [s_i, s_j]$ if $\varphi(s_i) \leq \varphi(s) \leq \varphi(s_j)$. Based on the string ordering, each B^{ed}-tree node N is associated with a string range $[N_b, N_e]$. Each leaf node contains f strings $\{s_1, \ldots, s_f\}$, where $s_i \in [N_b, N_e]$, for each $i \in [1, f]$. Each intermediate node contains multiple children nodes, where for each child of N, its range $[N'_b, N'_e] \subseteq [N_b, N_e]$, as $[N_b, N_e]$ is the range of N. We say these strings that stored in the sub-tree rooted at N as the strings that are *covered* by N.

We use $DST_{min}(s_q, N)$ to denote the minimal edit distance between a query string s_q and any s that is covered by a B^{ed}-tree node N. A nice property of the B^{ed}-tree is that, for any string s_q and node N, the string ordering φ enables to compute $DST_{min}(s_q, N)$ efficiently by computing $DST(s_q, N_b)$ and $DST(s_q, N_e)$ only, where N_b, N_e refer to the string range values of N. Then:

DEFINITION 3.1. Given a query string s_q and a similarity threshold θ, we say a B^{ed}-tree node N is a *candidate* if $DST_{min}(s_q, N) \leq \theta$. Otherwise, N is a *non-candidate*. ∎

B^{ed}-tree has an important monotone property. Given a node N_i in the B^{ed}-tree and any child node N_j of N_i, for any query string s_q, it must be true that $DST_{min}(s_q, N_i) \leq DST_{min}(s_q, N_j)$. Therefore, for any non-candidate node, all of its children must be non-candidates. This monotone property enables early termination of search on the branches that contain non-candidate nodes.

For any given query string s_q, the string similarity search algorithm starts from the root of the B^{ed}-tree, and iteratively visits the candidate nodes, until all candidate nodes are visited. The algorithm does not visit those non-candidate nodes as well as their descendants. An important note is that for each candidate node, some of its covered strings may still be dissimilar to the query string. Therefore, given a query string s_q and a candidate B^{ed}-tree node N that is associated with a range $[N_b, N_e]$, it is necessary to compute $DST(s_q, s)$, for each $s \in [N_b, N_e]$.

4. PROBLEM FORMULATION

In this section, we describe the authentication problem that we plan to study in this paper.

System Model. We consider the outsourcing model that involves three parties - a data owner who possesses a dataset D that contains n string values, the user (client) who requests for the similarity search on D, and a third-party service provider (server) that executes the similarity search on D. The data owner outsources D to the server. The server provides storage and similarity search as services. The client can be the data owner or a trusted party. Given the fact that the client may not possess D, we require that the availability of D is not necessary for the authentication procedure.

Similarity search queries. The server accepts similarity search queries from the clients. We consider the similarity search queries that take the format (S, θ), where S is a set of query strings, and θ is the similarity threshold. We consider two types of similarity search search queries: (1) *Single-string similarity search*: the query Q contains a single search string s_q (i.e., $S = \{s_q\}$). The server returns R that contains all similar strings of s_q in D; and (2) *Multi-string similarity search*: Q contains multiple unique search strings $S = \{s_1, \ldots, s_\ell\}$. The server returns the search results in the format of $\{\{s_1, R_{s_1}\}, \ldots, \{s_\ell, R_{s_\ell}\}\}$, where $R_{s_i} (1 \leq i \leq \ell)$ is the set of similar strings of s_i in D.

Note that the query strings may not necessarily exist in D. For each search string s_q, we consider two types of similarity outputs: (1) *Un-ranked results*: all the similar strings of s_q are returned without any ranking; and (2) *Top-k results*: the strings that are of the top-k smallest distances to s_q are returned by their distance to s_q in an ascending order. We consider both un-ranked and top-k ranking cases in the paper. We assume that the client sends a large number of similarity search queries to the server.

Our query model can be easily extended to support other types of string queries, e.g., range queries, KNN queries, and all-pairs join queries [33] that find all similar string pairs in two given string datasets.

Result correctness. We define the result correctness for two different types of query outputs.

Un-ranked search results. Given a query string s_q, the un-ranked search result R of s_q is *correct* if and only if it satisfies the following two conditions: (1) *soundness*: for any string $s' \in R$, s' must reside in D, and $s_q \approx s'$; and (2) *completeness*: for any string $s' \in D$ such that $s_q \approx s'$, s' must be included in R. In other words, for any string $s' \notin R$, it must be true that $s_q \not\approx s'$.

Top-k results. Given a query string s_q, let R be the ranked search result, in which the strings are ranked by their

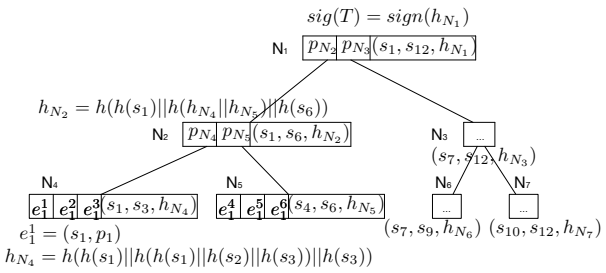

Figure 1: An example of MB-tree

distance to s_q in an ascending order. We use $R[i]$ $(1 \leq i \leq k)$ to denote the i-th string of R. Then R is *correct* if it satisfies the following two conditions: (1) *soundness*: $\forall 1 \leq i < j \leq k$, both $R[i]$ and $R[j]$ must exist in D, and $DST(s_q, R[i]) \leq DST(s_q, R[j])$; and (2) *completeness*: for any string $s' \notin R$, it must be true that $DST(s_q, s') \geq DST(s_q, R[k])$.

Threat model. In this paper, we assume that both the data owner and the client are fully trusted. However, the third-party server is not fully trusted as it could be compromised by the attacker (either inside or outside). The server may alter the received dataset D and return any search result that does not exist in D. It also may tamper with the search results. For instance, the server may return *incomplete* results that omit some legitimate documents in the similarity search results [24], or alter the ranking orders of the top-k results. Note that we do not consider privacy protection for the user queries. This issue can be addressed by private information retrieval (PIR) [4] and is beyond the scope of this paper.

5. SINGLE-STRING SIMILARITY SEARCH

In this section, we consider the single-string similarity search queries. We first present our basic verification approach named VS^2 (Section 5.1). Then we present our E-VS^2 method with improved verification cost (Section 5.2).

5.1 Basic Approach: VS^2

Given the dataset D, the query string s_q, and the distance threshold θ, the server returns all strings that are θ-similar to s_q. Besides the similar strings, the server also returns a *proof* of result correctness. In this section, we explain the details of our basic verification method of similarity search (VS^2). VS^2 consists of three phases: (1) the *pre-processing* phase in which the data owner constructs the authenticated data structure T of the dataset D. Both D and the root signature sig of T are outsourced to the server; (2) the *query processing* phase in which the server executes the similarity search query on D, and constructs the *verification object* (VO) of the search results R. The server returns both R and VO to the client; and (3) *verification* phase in which the client verifies the correctness of R by leveraging VO. Next we explain the details of these three phases.

5.1.1 Pre-Processing

In this one-time phase, the data owner constructs the authenticated data structure of the dataset D before outsourcing D to the server. We design a new authenticated data structure named the *Merkle B^{ed} tree* (MB-tree). Next, we explain the details of MB-tree.

The MB-tree is constructed on top of the B^{ed} tree by assigning the digests to each B^{ed} node. In particular, every MB-tree node contains a triple (N_b, N_e, h_N), where N_b, N_e correspond to the string range values associated with N, and h_N is the digest value computed as $h_N = h(h(N_b)||h(N_e)||h^{1\to f})$, where $h^{1\to f} = h(h_{C_1}||\ldots||h_{C_f})$, with $C_1, \ldots, C_f$ being the children of N. If N is a leaf node, then $C_1, \ldots, C_f$ are the strings $s_1, \ldots, s_f$ covered by N. Besides the triple, each MB-tree node contains multiple entries. In particular, for any leaf node N, assume it covers f strings. Then it contains f entries, each of the format (s, p), where s is a string that is covered by N, and p is the pointer to the disk block that stores s. For any intermediate node, assume that it has f children nodes. Then it contains f entries, each entry consisting of a pointer to one of its children nodes.

The digests of the MB-tree T can be constructed in the bottom-up fashion, starting from the leaf nodes. After all nodes of T are associated with the digest values, the data owner signs the root with her private key. The signature can be created by using a public-key cryptosystem (e.g., RSA). An example of the MB-tree structure is presented in Figure 1. The data owner sends both D and T to the server. The data owner keeps the root signature of T locally, and sends it to any client who requests for it for authentication purpose.

Following [5, 19], we assume that each node of the MB-tree occupies a disk page. For the constructed MB-tree T, each entry in the leaf node occupies $|s| + |p|$ space, where $|p|$ is the size of a pointer, and $|s|$ is the maximum length of a string value. The triple (N_b, N_e, h_N) takes the space of $2|s| + h$, where $|h|$ is the size of a hash value. Therefore, a leaf node can have $f_1 = [\frac{P-2|s|-|h|}{|p|+|s|}]$ entries at most, where P is the page size. Given n unique strings in the dataset, there are $[\frac{n}{f_1}]$ leaf nodes in T. Similarly, for the internal nodes, each entry takes the space of $|p|$. Thus each internal node can have at most $f_2 = [\frac{P-2|s|-|h|}{|p|}]$ entries (i.e., $[\frac{P-2|s|-|h|}{|p|}]$ children nodes). Therefore, the height h of T $h \geq log_{f_2}[\frac{n}{f_1}]$.

The construction complexity of MB-tree is $O(n)$, where n is the number of unique strings in D. It is cheaper than the complexity of pairwise similarity search over D.

5.1.2 VO Construction

Upon receiving the similarity search query $Q(s_q, \theta)$ from the client, the server calculates the edit distance between all the string pairs and distills the similar pairs. For the similarity search result R, the server constructs a verification object VO to show that R are both sound and complete.

First, we define *false hits*. Given a query string s_q and a similarity threshold θ, the *false hits* of s_q, denoted as F, are all the strings that are dissimilar to s_q. In other words, $F = \{s | s \in D, s_q \not\approx s\}$. Intuitively, to verify that R is sound and complete, the VO includes both similar strings R and false hits F. Apparently including all false hits may lead to a large VO, and thus high network communication cost and the verification cost at the client side. Therefore, we aim to reduce the VO size of F.

Before we explain how to reduce VO size, we first define C-strings and NC-strings. Apparently, each false hit string is covered by a leaf node of the MB-tree T. Based on whether a leaf node in MB-tree is a candidate, the false hits F are classified into two types: (1) **C-strings**: the strings that are covered by *candidate* leaf nodes; and (2) **NC-strings**: the strings that are covered by *non-candidate* leaf nodes.

Our key idea to reduce VO size of F is to include *representatives* of NC-strings instead of individual NC-strings. The representatives of NC-strings take the format of *maximal false hit subtrees* (MFs). Formally, given a MB-tree T, we say a subtree T^N that is rooted at node N is a *false hit*

subtree if N is a *non-candidate* node. We say the false hit subtree T^N rooted at N is *maximal* if the parent of N is a candidate node. The MFs can root at leaf nodes. Apparently, all strings covered by the MFs must be NC-strings. And each NC-string must be covered by a MF node. Furthermore, MFs are disjoint (i.e., no two MFs cover the same string). Therefore, instead of including individual NC-strings into the VO, we include their MFs. As the number of MFs is normally much smaller than the number of NC-strings, this can effectively reduce VO size. We are ready to define the VO.

DEFINITION 5.1. Given a dataset D, a query string s_q, let R be the returned similar strings of s_q. Let T be the MB-tree of D, and NC be the strings that are covered by non-candidate nodes of T. Let $\mathcal{M}$ be a set of MFs of NC. Then the VO of s_q consists of: (i) string s, for each $s \in D - NC$; and (ii) a pair $(N, h^{1\to f})$ for each $MF \in \mathcal{M}$ that is rooted at node N, where N is represented as $[N_b, N_e]$, with $[N_b, N_e]$ the string range associated with N, and $h^{1\to f} = h(h_{C_1}||\ldots||h_{C_f})$, with $C_1, \ldots, C_f$ being the children of N. If N is a leaf node, then $C_1, \ldots, C_f$ are the strings $s_1, \ldots, s_f$ covered by N. Furthermore, in VO, a pair of brackets is added around the strings and/or the pairs that share the same parent in T. ∎

Intuitively, in VO, the similar strings and C-strings are present in the original string format, while NC-strings are represented by the MFs (i.e., in the format of $([N_b, N_e], h_N)$).

EXAMPLE 5.1. Consider the MB-tree T in Figure 1, and the query string s_1. Note that $s_1 \in D$ (but in general s_q may not be present in D). Assume the similar strings are $R = \{s_1, s_3, s_5\}$. Also assume that node N_6 of T is the only non-candidate node. Then NC-strings are $NC = \{s_7, s_8, s_9\}$, and C-strings are $\{s_2, s_4, s_6, s_{10}, s_{11}, s_{12}\}$. The set of MFs $\mathcal{M} = \{N_6\}$. Therefore,

$$VO(s_1) = \{(((s_1, s_2, s_3), (s_4, s_5, s_6)), (([s_7, s_9], h^{7\to 9}), (s_{10}, s_{11}, s_{12})))\}, where\ h^{7\to 9} = h(h(s_7)||h(s_8)||h(s_9))$$ ∎.

For each MF, we do not require that $h(N_b)$ and $h(N_e)$ appear in VO. However, the server must include both N_b and N_e in VO. This is to prevent the server to cheat on the non-candidate nodes by including incorrect $[N_b, N_e]$ in VO. More details of the robustness of our authentication procedure can be found in Section 5.1.4.

5.1.3 Authentication Phase

For a given query string s_q, the server returns $\{R, VO\}$ to the client. Before result authentication, the client obtains the data owner's public key from a certificate authority (e.g., VeriSign [14]). The client also obtains the hash function, the root signature sig of the MB-tree, and the string ordering scheme and the B^{ed}-tree construction procedure from the data owner. The verification procedure consists of three steps. In Step 1, the client re-constructs the MB-tree from VO. In Step 2, the client re-computes the root signature sig', and compares sig' with sig. In Step 3, the client re-computes the edit distance between s_q and a *subset* of strings in VO. Next, we explain the details of these steps.

Step 1: Re-construction of MB-tree: First, the client sorts the strings and string ranges (in the format of $[N_b, N_e]$) in VO by their mapping values according to the string ordering scheme. String s is put ahead the range $[N_b, N_e]$ if $s < N_b$. It should return a total order of strings and string ranges. If there exists any two ranges $[N_b, N_e]$ and $[N'_b, N'_e]$ that overlap, the client concludes that the VO is not correct. If there exists a string $s \in R$ and a range $[N_b, N_e] \in VO$ such that $s \in [N_b, N_e]$, the client concludes that R is not sound, as s indeed is a dissimilar string (i.e., it is included in a non-candidate node). Second, the client maps each string $s \in R$ to an entry in a leaf node in T, and each pair $([N_b, N_e], h_N) \in VO$ to an internal node in T. The client re-constructs the parent-children relationships between these nodes by following the matching brackets () in VO.

Step 2: Re-computation of root signature: After the MB-tree T is re-constructed, the client computes the root signature of T. For each string value s, the client calculates $h(s)$, where $h()$ is the same hash function used by the data owner for the construction of the MB^{ed}-tree. For each internal node that corresponds to a pair $([N_b, N_e], h^{1\to f})$ in VO, the client computes the hash h_N of N as $h_N = h(h(N_b)||h(N_e)||h^{1\to f})$. Finally, the client re-computes the hash value of the root node, and rebuilds the root signature sig' by signing h_{root} using the data owner's public key. The client then compares sig' with sig. If $sig' \neq sig$, the client concludes that the server's results are not correct.

Step 3: Re-computation of necessary edit distance: First, for each string $s \in R$, the client re-computes the edit distance $DST(s_q, s)$, and verifies whether $DST(s_q, s) \leq \theta$. If all strings $s \in R$ pass the verification, then the client concludes that R is *sound*. Second, for each C-string $s \in VO$ (i.e., those strings appear in VO but not R), the client verifies whether $DST(s_q, s) > \theta$. If it is not (i.e., s is a similar string indeed), the client concludes that the server fails the completeness verification. Third, for each range $[N_b, N_e] \in VO$, the client verifies whether $DST_{min}(s_q, N) > \theta$, where N is the corresponding MB-tree node associated with the range $[N_b, N_e]$. If it is not (i.e., node N is indeed a candidate node), the client concludes that the server fails the soundness verification.

EXAMPLE 5.2. Consider the MB-tree in Figure 1 as an example, and the query string s_1. Assume the similar strings $R = \{s_1, s_3, s_5\}$. Consider the VO shown in Example 5.1. The C-strings are $C = \{s_2, s_4, s_6, s_{10}, s_{11}, s_{12}\}$. After the client re-constructs the MB-tree, it re-computes the hash values of strings $R \cup C$. It also computes the digest $h_{N_6} = (h(s_7)||h(s_9)||h^{7\to 9})$. Then it computes the root signature sig' from these hash values. It also performs the following distance computations: (1) for $R = \{s_1, s_3, s_5\}$, compute the edit distance between s_1 and string in R, (2) for $C = \{s_2, s_4, s_6, s_{10}, s_{11}, s_{12}\}$, compute the edit distance between s_1 and any C-string in C, and (3) for the pair $([s_7, s_9], h_{N_6}) \in VO$, compute $DST_{min}(s_q, N_6)$. ∎

5.1.4 Security Analysis

Given a query string s_q and a similarity threshold θ, let R (F, resp.) be the similar strings (false hits, resp.) of s_q. An untrusted server may perform the following cheating behaviors the real results R: (1) *tampered values*: some strings in R do no exist in the original dataset D; (2) *soundness violation*: the server returns $R' = R \cup FS$, where $FS \subseteq F$; and (3) *completeness violation*: the server returns $R' = R - SS$, where $SS \subseteq R$.

The tampered values can be easily caught by the authentication procedure, as the hash values of the tampered strings are not the same as the original strings. This leads to that the root signature of MB-tree re-constructed by the client

differet from the root signature from the data owner. The client can catch the tampered values by Step 2 of the authentication procedure. Next, we mainly focus on the discussion of how to catch soundness and completeness violations.

Soundness. The server may deal with the VO construction of $R' = R \cup FS$ in two different ways:

Case 1. The server constructs the VO V of the correct result R, and returns $\{R', V\}$ to the client;

Case 2. the server constructs the VO V' of R', and returns $\{R', V'\}$ to the client. Note that the strings in FS can be either NC-strings or C-strings. Next, we discuss how to catch these two types of strings for both Case 1 and 2.

For Case 1, for each NC-string $s \in FS$, s must fall into a MF-tree node in V. Thus, there must exist an MF-tree node whose associated string range overlaps with s. The client can catch s by Step 1 of the authentication procedure. For each C-string $s \in FS$, s must be treated as a C-string in V. Thus the client can catch it by computing $DST(s, s_q)$ (i.e., Step 3 of the authentication procedure).

For Case 2, the C-strings in FS will be caught in the same way as Case 1. Regarding NC-strings in FS, they will not be included in any MF-tree in V'. Therefore, the client cannot catch them by Step 1 of the authentication procedure (as Case 1). However, as these strings are included in R', the client still can catch these strings by computing the edit distance of s_q and any string in R' (i.e., Step 3 of the authentication procedure).

Completeness. To deal with the VO construction of $R' = R - SS$, we again consider the two cases as for the discussion of correctness violation. In particular, let V and V' be the VO constructed from the correct result R and the incomplete result $R' = R - SS$ respectively, where $SS \subseteq R$. We discuss how to catch these two cases in details.

For Case 1 (i.e., the server returns $\{R', V\}$), any string $s \in SS$ is a C-string. These strings can be caught by re-computing the edit distance between the query string and any C-string (i.e., Step 3 of the authentication procedure).

For Case 2 (i.e., the server returns $\{R', V'\}$), any string $s \in SS$ is either a NC-string or a C-string in V'. For any C-string $s \in SS$, it can be caught by re-computing the edit distance between the query string and the C-strings (i.e., Step 3 of the authentication procedure). For any NC-string $s \in SS$, it must be included into a non-candidate MB-tree node. We have the following theorem.

THEOREM 5.1. Given a query string s_q and a non-candidate node N of range $[N_b, N_e]$, including any string s' into N such that $s' \approx s_q$ must change N to be a candidate node.

The proof is straightforward. It is easy to see that $s' \notin [N_b, N_e]$. Therefore, including s' into N must change the range to be either $[s', N_e]$ or $[N_b, s']$, depending on whether $s' < N_b$ or $s' > N_e$. Now it must be true that $DST_{min}(s_q, N) \leq \theta$, as $DST(s', s_q) \leq \theta$.

Following Theorem 5.1, for any string $s \in SS$ that is considered a NC-string in V', the client can easily catch it by verifying whether $DST_{min}(s_q, N) > \theta$, for any non-candidate node (i.e., Step 3 of the authentication procedure).

5.2 String Embedding Authentication: E-VS^2

One weakness of VS^2 is that if there exist a significant number of C-strings, its VO can be of large size. This may bring expensive network communication cost. Furthermore, since the client needs to compute the edit distance $DST(s_q, s)$ for each C-string s, too many C-strings may incur expensive verification cost at the client side too. Our goal is to shrink VO size with regard to C-strings, so that both network communication cost and verification cost can be reduced. We observe that although C-strings are not similar to the query string, they may be similar to each other. Therefore, we design *E-VS^2*, a computation-efficient method on top of VS^2. The key idea of E-VS^2 is to construct a set of *representatives* of C-strings based on their similarity, and only include the representatives of C-strings in VO. To construct the representatives of C-strings, we first apply a similarity-preserving mapping function on C-strings, and transform them into the Euclidean space, so that the similar strings are mapped to the close points in the Euclidean space. Then C-strings are organized into a small number of groups called *distant bounding hyper-rectangles (DBHs)*. DBHs are the representatives of C-strings in VO. In the verification phase, the client only needs to calculate the Euclidean distance between s_q and DBHs. Since the number of DBHs is much smaller than the number of C-strings, and Euclidean distance calculation is much faster than that of edit distance, the verification cost of the E-VS^2 approach is much cheaper than that of VS^2. Next, we explain the details of the E-VS^2 approach. Similar to the VS^2 approach, E-VS^2 consists of three phases: (1) the *pre-processing* phase at the data owner side; (2) the *query processing* phase at the server side; and (3) the *verification* phase at the client side. Next, we discuss the three phases in details.

5.2.1 Pre-Processing

Before outsourcing the dataset D to the server, similar to VS^2, the data owner constructs the MB-tree T on D. In addition, the data owner maps D to the Euclidean space E via a similarity-preserving embedding function $f : D \rightarrow E$ denote the embedding function. We use SparseMap [13] as the embedding function due to its contractive property (Section 3). The complexity of the embedding is $O(cdn^2)$, where c is a constant value between 0 and 1, d is the number of dimensions of the Euclidean space, and n is the number of strings of D. We agree that the complexity of string embedding is comparable to the complexity of similarity search over D. This naturally fits into the amortized model for outsourced computation [9]: the data owner performs a one-time computationally expensive phase (in our case constructing the embedding space), whose cost is amortized over the authentication of all the future query executions.

The data owner sends D, T and the embedding function f to the server. The server constructs the embedded space of D by using the embedding function f. The function f will also be available to the client for result authentication.

5.2.2 VO Construction

Given the query $Q(s_q, \theta)$ from the client, the server applies the embedding function f on s_q, and finds its corresponding node P_q in the Euclidean space. Then the server finds the result set R of s_q. To prove the soundness and completeness of R, the server builds a verification object VO. First, similar to VS^2, the server searches the MB-tree to build MFs of NC-strings. For the C-strings, the server constructs a set of *distant bounding hyper-rectangles* (DBHs) from their embedded nodes in the Euclidean space. Before we define DBH, first, we define the minimum distance between an Euclidean point and a hyper-rectangle. Given a set of points $\mathcal{P}$ $\{P_1, \ldots, P_t\}$ in a d-dimensional Euclidean space, a hyper-rectangle $R(< l_1, u_1 >, \ldots, < l_d, u_d >)$ is the *minimum bounding hyper-rectangle* (MBH) of $\mathcal{P}$ if $l_i = min_{k=1}^{t}(P_k[i])$

and $u_i = max_{k=1}^{t}(P_k[i])$, for $1 \leq i \leq d$, where $P_k[i]$ is the i-dimensional value of P_k. For any point P and any hyper-rectangle $R(< l_1, u_1 >, \ldots, < l_d, u_d >)$, the minimum Euclidean distance between P and R is $dst_{min}(P, R) = \sqrt{\sum_{1 \leq i \leq d} m[i]^2}$, where $m[i] = max\{l_i - p[i], 0, p[i] - u_i$. Intuitively, if the node P is inside R, the minimum distance between P and R is 0. Otherwise, we pick the length of the shortest path that starts from P to reach R. We have:

LEMMA 5.1. Given a point P and a hyper-rectangle R, for any point $P' \in R$, the Euclidean distance $dst(P', P) > dst_{min}(P, R)$.

The proof of Lemma 5.1 is trivial. We omit the details due to the space limit.

Now we are ready to define *distant bounding hyper-rectangles* (DBHs). Given a query string s_q, let P_q be its embedded point in the Euclidean space. For any hyper-rectangle R in the same space, R is a *distant bounding hyper-rectangle* (DBH) of P_q if $dst_{min}(P_q, R) > \theta$.

Given a DBH R, Lemma 5.1 guarantees that $dst(P_q, P) > \theta$ for any point $P \in R$. Recalling the contractive property of the SparseMap method, we have $dst(P_i, P_j) \leq DST(s_i, s_j)$ for any string pair s_i, s_j and their embedded points P_i and P_j. Thus we have the following theorem:

THEOREM 5.2. Given a query string s_q, let P_q be its embedded point. Then for any string s, s must be dissimilar to s_q if there exists a DBH R of P_q such that $P \in R$, where P is the embedded point of s.

Based on Theorem 5.2, to prove that the C-strings are dissimilar to the query string s_q, the server can build a number of DBHs to the embedded Euclidean points of these C-strings. We must note that not all C-strings can be included into DBHs. This is because the embedding function may introduce *false positives*, i.e., there may exist a false hit string s of s_q whose embedded point P becomes close to P_q. Given a query string s_q, we say a C-string s of s_q is an *FP-string* if $dst(P, P_q) \leq \theta$, where P and P_q are the embedded Euclidean points of s and s_q. Otherwise (i.e. $dst(P, P_q) > \theta$), we call s a *DBH-string*. We have:

THEOREM 5.3. Given a query string s_q, for any DBH-string s, its embedded point P must belong to a DBH.

The proof of Theorem 5.3 is straightforward. For any DBH-string whose embedded point cannot be included into a DBH with other points, it constructs a hyper-rectangle H that only consists of one point. Obviously H is a DBH.

Therefore, given a query string s_q and a set of C-strings C, first, the server classifies C into FP- and DBH-strings, based on the Euclidean distance between their embedded points. Apparently the embedded points of FP-strings cannot be put into any DBH. Therefore, the server only consider DBH-strings and tries builds DBHs of DBH-strings. The VO of DBH-strings will be computed from DBHs. Therefore, in order to minimize the verification cost at the client side, the server aims to minimize the number of DBHs. Formally,

$MDBH$ Problem: Given a set of DBH-strings $\{s_1, \ldots, s_t\}$, let $\mathcal{P}$ $\{P_1, \ldots, P_t\}$ be their embedded points. Construct a minimum number of DBHs $\mathcal{R} = \{R_1, \ldots, R_k\}$ such that: (1) $\forall R_i, R_j \in \mathcal{R}$, R_i and R_j do not overlap; and (2) $\forall P_i \in \mathcal{P}$, there exists a DBH $R \in \mathcal{R}$ such that $P_i \in R$.

Next, we present the solution to the $MDBH$ problem. We first present the simple case for a 2-D dimension (i.e. $d = 2$). Then we discuss the scenario when $d > 2$. For both settings, consider the same input that includes a query point P_q and set of Euclidean points $\mathcal{P}$ $\{P_1, \ldots, P_t\}$ which are the embedded points of a query string s_q and its DBH-strings respectively.

When $d = 2$. We construct a graph $G = (V, E)$ such that for each point $P_i \in \mathcal{P}$, it corresponds to a vertex $v_i \in V$. For any two vertices v_i and v_j that correspond to two points P_i and P_j, there is an edge $(v_i, v_j) \in E$ if $dst_{min}(P_q, R) > \theta$, where R is the MBH of P_i and P_j. We have:

THEOREM 5.4. Given the graph $G = (V, E)$ constructed as above, for any clique C in G, let R be the MBH constructed from the points corresponding to the vertice in C. Then R must be a DBH.

The proof of Theorem 5.4 is in Appendix. Based on Theorem 5.4, the *MDBH* problem is equivalent to the well-known *clique partition problem*, which is to find the smallest number of cliques in a graph such that every vertex in the graph belongs to exactly one clique. The clique partition problem is NP-complete. Thus, we design the heuristic solution to our *MDBH* problem. Our heuristic algorithm is based on the concept of *maximal cliques*. Formally, a clique is maximal if it cannot include one more adjacent vertex. The maximal cliques can be constructed in polynomial time [7]. It is shown that every maximal clique is part of some optimal clique-partition [7]. Based on this, finding a minimal number of cliques is equivalent to finding a number of maximal cliques. Thus we construct maximal cliques of G iteratively, until all the vertices belong to at least one clique.

There is a special case where the $MDBH$ problem can be solved in polynomial time: when the embedded points of all DBH-strings lie on a single line, we can construct a minimal number of DBHs in the complexity of $O(\ell)$, where ℓ is the number of DBH-strings. Due to space limit, the details of DBH construction can be found in Appendix.

When $d > 2$. Unfortunately, Theorem 5.4 can not be extended to the case of $d > 2$. We found an example in which the MBHs of the pairs (v_i, v_j), (v_i, v_k), and (v_j, v_k) are DBHs. However, the MBH of the pair (v_i, v_j, v_k) is not a DBH, as it includes a point w such that w is not inside $R(v_i, v_j)$, $R(v_i, v_k)$, and $R(v_j, v_k)$, but $dst(P_q, w) < \theta$.

To construct the DBHs for the case $d > 2$, we slightly modify the clique-based construction algorithm for the case $d = 2$. In particular, when we extend a clique C by adding an adjacent vertex v, we check if the MBH of the extended clique $C' = C \cup \{v\}$ is a DBH. If not, we delete the edges (u, v) from G for all $u \in C$. This step ensures that if we merge any vertex in U_1 to C, the MBH of the newly generated clique is still a DBH.

For both cases $d = 2$ and $d > 2$, the complexity of constructing DBHs from DBH-strings is $O(n_{DS}^3)$, where n_{DS} is the number of DBH-strings.

Now we are ready to describe VO construction by the E-VS^2 approach. Given a dataset D and a query string s_q, let R and F be the similar strings and false hits of s_q respectively. VS^2 approach groups F into C-strings and NC-strings. E-VS^2 approach further groups C-strings into FP-strings and DBH-strings. Then E-VS^2 constructs VO from R, NC-strings, FP-strings, and DBH-strings. Formally,

DEFINITION 5.2. Given a query string s_q, let R be the returned similar strings of s_q. Let NC be the NC-strings, and DS the DBH-strings. Let T be the MB-tree, MF be the maximum false hit trees of NC. Let $\mathcal{R}$ be the set of DBH constructed from DBH. Then the VO of s_q consists of: (i)

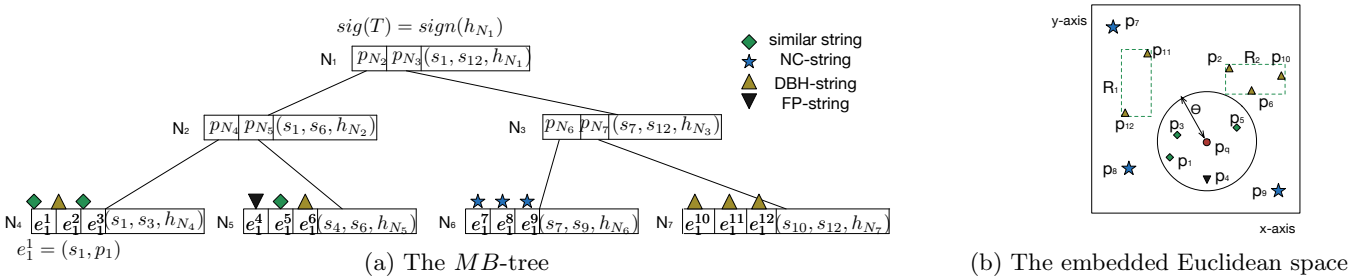

(a) The MB-tree (b) The embedded Euclidean space

Figure 2: An example of VO construction by E-VS^2 method

string s, for each $s \in D - NC - DS$; (ii) a pair $(N, h^{1\to f})$ for each non-leaf MF that is rooted at node N, where N takes the format of $[N_b, N_e]$, with $[N_b, N_e]$ the string range associated with N, and $h^{1\to f} = h(h_{C_1}||\ldots||h_{C_f})$, with $C_1, \ldots, C_f$ being the children of N; (iii) $\mathcal{R}$; and (iv) a pair (s, p_R) for each $s \in DS$, where p_R is the pointer to the DBH in $\mathcal{R}$ that covers the Euclidean point of s; Furthermore, in VO, a pair of square bracket is added around the strings pairs that share the same parent in T. ∎

Example 5.3. To continue with our running example in Example 5.2, recall that the query string is s_1. The similar strings $R = \{s_1, s_3, s_5\}$. The NC-strings $NC = \{s_7, s_8, s_9\}$. The C-strings $C = \{s_2, s_4, s_6, s_{10}, s_{11}, s_{12}\}$. Consider the embedded Euclidean space shown in Figure 2 (b). Apparently s_4 is a FP-string as $dst(p_q, p_4) < \theta$. So the DBH-strings are $\{s_2, s_6, s_{10}, s_{11}, s_{12}\}$. The DBHs of these DBH-strings are shown in the rectangles in Figure 2 (b). The VO of query string s_q is

$$VO(s_q) = \{(((s_1, (s_2, p_{R_2}), s_3), (s_4, s_5, (s_6, p_{R_2}))), ([s_7, s_9], h^{7\to 9}), ((s_{10}, p_{R_2}), (s_{11}, p_{R_1}), (s_{12}, p_{R_1}))), \{R_1, R_2\}\},$$

where $h^{7\to 9} = h(h(s_7)||h(s_8)||h(s_9)))$. ∎

5.2.3 VO-based Authentication

After receiving (R, VO) from the server, the client uses VO to verify if R is sound and complete. The verification of E-VS^2 consists of four steps. The first three steps are similar to the three steps of the VS^2 approach. The fourth step is to re-compute a set of Euclidean distance. Next, we discuss the four steps in details.

Step 1 & 2: these two steps are exactly the same as Step 1 & 2 of VS^2.

Step 3: Re-computing necessary edit distance: Similar to VS^2, first, for each $s \in R$, the client verifies $DST(s, s_q) \leq \theta$. Second, for each range $[N_b, N_e] \in VO$, the client verifies whether $DST_{min}(s_q, N) > \theta$, where N is the corresponding MB-tree node associated with the range $[N_b, N_e]$. The only difference of the E-VS^2 approach is that for each FP-string s, the client verifies if $DST(s_q, s) > \theta$. If not, the client concludes that R is incomplete.

Step 4: Re-computing of necessary Euclidean distance: Step 3 only verifies the dissimilarity of FP- and NC-strings. In this step, the client verifies the dissimilarity of DBH-strings. First, for each pair $(s, p_R) \in VO$, the client checks if $P_s \in R$, where P_s is the embedded point of s, and R is the DBH that p_R points to. If all pairs pass the verification, the client ensures that the DBHs in VO covers the embedded points of all the DBH-strings. Second, for each DBH $R \in VO$, the client checks if $dst_{min}(P_q, R) > \theta$. If it is not, the client concludes that the returned results are not correct. Otherwise, third, for each similar string $s \in R$, the client checks if there exists any DBH that includes P_s, where P_s is the embedded point of s. If there does, the client concludes that the results violate soundness.

Note that we do not require to re-compute the edit distance between any DBH-string and the query string. Instead we only require the computation of the Euclidean distance between a set of DBHs and the embedded points of the query string. Since Euclidean computation is much faster than that of the edit distance. Therefore E-VS^2 saves much verification cost compared with VS^2. More comparison of VS^2 and E-VS^2 can be found in Section 5.3.1.

Example 5.4. Following the running example in Example 5.3, after calculating the root signature sig' from VO and compares it with the signature sig received from the data owner, the client performs the following computations: (1) for $R = \{s_1, s_3, s_5\}$, compute the edit distances between s_1 and any string in R; (2) for NC-strings $NC = \{s_7, s_8, s_9\}$, compute $DST_{min}(s_q, N_3)$; (3) for FP-strings $FP = \{s_4\}$, compute the edit distance $DST(s_q, s_4)$; and (4) for DBR-strings $DS = \{s_2, s_6, s_{10}, s_{11}, s_{12}\}$, compute $dst_{min}(P_q, R_1)$ and $dst_{min}(P_q, R_2)$. Compared with the VS^2 approach in Example 5.2, which computes 9 edit distances, E-VS^2 computes 4 edit distances, and 2 Euclidean distances. Recall that the computation of Euclidean distance is much cheaper than that of the edit distance. ∎

5.3 Security Analysis

Similar to the security discussion for the VS^2 approach (Sec. 5.1.4), the server may perform three types of cheating behaviors, i.e., tampered values, soundness violation, and completeness violation. E-VS^2 can catch the cheating behaviors of tampered values by re-computing the root signature of MB-tree (i.e., Step 2 of the authentication procedure). Next, we mainly focus on how to catch the correctness and completeness violation by the E-VS^2 approach.

Soundness. To violate soundess, the server returns $R' = R \cup FS$, where $FS \subseteq F$ (i.e., FS is a subset of false hits). We consider two possible ways that the server constructs VO:

Phase	Measurement	VS^2	E-VS^2
Pre-processing	Time	$O(n)$	$O(cdn^2)$
	Space	$O(n)$	$O(n)$
VO construction	time	$O(n)$	$O(n + n_{DS}^3)$
	VO Size	$(n_R + n_C)\sigma_S + n_{MF}\sigma_M$	$(n_R + n_F)\sigma_S + n_{MF}\sigma_M + n_{DBH}\sigma_D$
Verification	Time	$O((n_R + n_{MF} + n_C)C_{Ed})$	$O((n_R + n_{MF} + n_F)C_{Ed} + n_{DBH}C_{El})$

Table 1: Complexity comparison between VS^2 and E-VS^2
(n: # of strings in D; c: a constant in [0, 1]; d: # of dimensions of Euclidean space; σ_S: the average length of the string; σ_M: Avg. size of a MB-tree node; σ_D: Avg. size of a DBH; n_R: # of strings in R; n_C: # of C-strings; n_F: # of FP-strings; n_{DS}: # of DBH-strings; n_{DBH}: # of DBHs; n_{MF}: # of MF nodes; C_{Ed}: the complexity of an edit distance computation; C_{El}: the complexity of Euclidean distance calculation.)

(Case 1.) the server constructs the VO V of the correct result R, and returns $\{R', V\}$ to the client; and Case 2.) the server constructs the VO V' of R', and returns $\{R', V'\}$ to the client. Note that the strings in FS can be NC-strings, FP-strings, and DBH-strings. Next, we discuss how to catch these three types of strings for both cases.

For Case 1, for any NC-string $s \in FS$, s can be caught in the same way as by VS^2 approach (i.e., Step 1 of authentication). For any FP-string $s \in FS$, s can be caught by re-computing the edit distance $DST(s, s_q)$ (i.e., Step 3 of authentication). For any DBH-string $s \in FS$, let P_s be the embedded point of s. Since the VO V is constructed from the correct R and F, there must exist a DBH in V that includes P_s. Therefore, s can be caught by verifying whether there exist any DBH that includes the embedded points of s (Step 4 of the authentication procedure).

For Case 2, the NC-strings and FP-strings in FS can be caught in the same way as in Case 1. The DBH-strings cannot be caught by Step 4 now, as: (1) the DBHs constructed from a subset of DBH-strings are still DBHs, and (2) no string in FS is included in a DBH in the VO V'. However, these DBH-strings are treated as FP-strings (i.e., not included in any DBH), and thus can be caught by Step 3.

Completeness. To violate the completeness requirements, the server returns returns $R' = R - SS$, where $SS \subseteq R$. Let V and V' be the VO constructed from R and R' respectively. We again consider the two cases as for the discussion of soundness violation.

For Case 1 (i.e., the server returns $\{R', V\}$), any string $s \in SS$ is a FP-string, as s is not included in R. Then by calculating $DST(s, s_q)$, (i.e., Step 3 of the authentication procedure), the client discovers that s is indeed similar to s_q and thus catch the incomplete results.

For Case 2 (i.e, the server returns $\{R', V'\}$), any string $s \in SS$ is either a FP-string or a DBH-string. If s is treated as a FP-string, it can be caught by recomputing of edit distance (Step 3 of the authentication procedure). If s is a DBH-string, then its Euclidean point P_s must be included in a DBH R. We have the following theorem.

THEOREM 5.5. Given a query string s_q and a DBH R (i.e., $dst_{min}(P_q, R) > \theta$), then for any string s such that $s \approx s_q$, adding the embedded point of s to R must change R to be a non-DBH.

The proof of Theorem 5.5 is straightforward. Let R' be the hyper-rectangle after adding P_s to R. It must be true that $dst_{min}(P_q, R) \leq dst(P_q, P) \leq dst(s_q, s) \leq \theta$. Then R' cannot be a DBH.

Following Theorem 5.5, for any string $s \in SS$, including its embedded point into any DBH R will change R to be a non-DBH. Then the client can easily catch it by re-computing the euclidean distance (Step 4 of verification).

5.3.1 VS^2 Versus E-VS^2

In this section, we compare VS^2 and E-VS^2 approaches in terms of the time and space of the pre-processing, VO construction, and verification phases. The comparison results are summarized in Table 1. Regarding the VO construction overhead at the server side, as shown in our empirical study, $n_{DS} << n$, thus the overhead $O(n + n_{DS}^3)$ of the E-VS^2 approach is comparable to $O(n)$ of the VS^2 approach.

Regarding the VO size, the VO size of the VS^2 approach is calculated as the sum of two parts: (1) the total size of the similar strings and C-strings (in string format), and (2) the size of MF nodes. Note that $\sigma_M = 2\sigma_S + |h|$, where $|h|$ is the size of a hash value. In our experiments, it turned out that $\sigma_M/\sigma_S \approx 10$. The VO size of the E-$VS^2$ approach is calculated as the sum of three parts: (1) the total size of the similar strings and FP-strings (in string format), (2) the size of MF nodes, and (3) the size of DBHs. Our experimental results show that $\sigma_D >> \sigma_S, \sigma_M$.

Regarding the complexity of verification time, note that $n_C = n_F + n_{DS}$, where n_C, n_F and n_{DS} are the number of C-strings, FP-strings, and DBH-strings respectively. Usually, $n_{DBH} < n_{DS}$ as a single DBH can cover the Euclidean points of a large number of DBH-strings. Also note that C_{Ed} (i.e., complexity of an edit distance computation) is much more expensive than C_{El} (i.e, the complexity of Euclidean distance calculation). Our experiments show that the time to compute one single edit distance can be 20 times of computing one Euclidean distance. Therefore, compared with VS^2, E-VS^2 significantly reduces the verification overhead at the client side. We admit that it increases the overhead of pre-processing at the data owner side and the VO construction at the server side. We argue that, as the pre-processing phase is a one-time operation, the cost of constructing the embedding function can be amortized by a large number of queries from the client.

6. MULTI-STRING SIMILARITY SEARCH

So far we discussed the authentication of single-string similarity search queries. To authenticate a multi-string query $Q(S, \theta)$ that contains multiple unique search strings $S = \{s_1, \ldots, s_\ell\}$, a straightforward solution is to create VO for each string $s_i \in S$ and its similarity result R_i. Apparently this solution may lead to VO of large sizes in total. Thus, we aim to reduce the size of VOs. Our VO optimization method consists of two main strategies: (1) optimization by triangle inequality; and (2) optimization by overlapping dissimilar strings.

6.1 Optimization by Triangle Inequality

It is well known that the string edit distance satisfies the *triangle inequality*, i.e. $|DST(s_i, s_k) - DST(s_j, s_k)| \leq$

$DST(s_i, s_j) \leq DST(s_i, s_k) + DST(s_j, s_k)$. Therefore, consider two query strings s_{q_1} and s_{q_2}, assume the the server has executed the similarity search of s_{q_1} and prepared the VO of the results. Then consider the authentication of the search results of string s_{q_2}, for any string $s \in D$ such that $DST(s, s_{q_1}) - DST(s_{q_1}, s_{q_2}) > \theta$, there is no need to prepare the proof of s showing that it is a false hit for s_{q_2}. A straightforward method is to remove s from the MB-tree T for VO construction for query string s_{q_2}. Now the question is whether removing s always lead to VO of smaller sizes. We have the following theorem.

THEOREM 6.1. Given a string s and a MB-tree T, let $N \in T$ be the corresponding node of s, and N_P be the parent of N in T. Let $C(N_P)$ be the set of children of P in T, and N'_P be the node constructed from $C(N_P) \backslash N$ (i.e., the children nodes of N_p excluding N). Then if N_P is a non-candidate, it must be true that N'_P must be a non-candidate.

Let $[N_b, N_e]$ be the range of N_P. The proof of Theorem 6.1 considers two possible cases: (1) $s \neq N_b$ and $s \neq N_e$. Then removing s will not change N_b and N_e, which results in that both N'_p and N_p have the same range $[N_b, N_e]$. (2) $s = N_b$ or $s = N_e$. Then removing s from $C(N_P)$ changes the range of N_P to be $[N'_b, N'_e]$. Note that it must be true $[N'_b, N'_e] \subseteq [N_b, N_e]$. Therefore, N'_P must be a non-candidate.

Based on Theorem 6.1, removing the MB-tree nodes that correspond to any dissimilar string indeed does not change the structure of the MB-tree. Therefore, we can optimize the VO construction procedure by removing those strings covered by the triangle inequality from the MB-tree.

Another possible optimization by triangle inequality is that for any two query strings s_{q_1} and s_{q_2}, for any string $s \in D$ such that $DST(s, s_{q_1}) + DST(s_{q_1}, s_{q_2}) \leq \theta$, then the client does not need to re-compute $DST(s, s_{q_2})$ proving that s_j is a similar string in the results. We omit the details due to the space limit.

6.2 Optimization by Overlapped Dissimilar Strings

Given multiple-string search query, the key optimization idea is to merge the VOs of individual query strings. This is motivated by the fact that any two query strings s_i and s_j may share a number of dissimilar strings. These shared dissimilar strings can enable to merge the VOs of s_i and s_j. Note that simply merging all similar strings of s_i and s_j into one set and constructing MB-tree of the merged set is not correct, as the resulting MB-tree may deliver non-leaf MFs that are candidates to both s_i and s_j. Therefore, given two query strings s_i and s_j such that their false hits overlap, let NC_i (NC_j, resp.) and DBH_i (DBH_j, resp.) be the NC-strings and DBH-strings of s_i (s_j resp.), the server finds the overlap of NC_i and NC_j, as well as the overlap between DBH_i and DBH_j. Then the server constructs VO that shares the same data structure on these overlapping strings. In particular, first, given the overlap $O_1 = NC_i \cap NC_j$, the server constructs the non-leaf MFs from O_1, and include the constructed MFs in the VOs of both s_i and s_j. Second, given the overlap $O_2 = DBH_i \cap DBH_j$, the server constructs the $DBHs$ from O_2, and include the constructed $DBHs$ in the VOs of both s_i and s_j.

7. AUTHENTICATION OF TOP-K RESULTS

In this section, we discuss how to authenticate the top-k similarity results. Formally, given a query string s_q and a threshold value $\theta > 0$, the query is to find the top-k similar strings of s, which are sorted by their distance to the query string s_q in ascending order. In other words, the query returns a set of strings $R = \{s_1, s_2, \ldots, s_k\}$, where $DST(s, s_i) \leq \theta$ and $DST(s, s_i) \leq DST(s, s_{i+1}), \forall i \in [1, k-1]$. Let c be the number of strings that are similar to s in terms of θ. We consider two cases: (1) $k = c$ (i.e., all similar strings are ranked); and (2) $k < c$ (i.e., only a subset of similar strings are returned).

k=c. Besides verifying the soundness and completeness of the returned strings, the authentication procedure checks if the ranking is correct. Both the soundness and completeness verification can be achieved by our VS^2 or E-VS^2 approach. A straightforward solution to ranking authentication is to calculate the pairwise distance $DST(s_q, s)$ for any $s \in R$, re-sort strings in R based on their distances, and compare the ranking results with R. Since the calculation of pairwise distance is required for soundness verification (for both VS^2 and E-VS^2), the authentication can be done by one additional sorting step.

k<c. Given a query string s_q and the top-k similar strings $R = \{s_1, s_2, \ldots, s_k\}$, the client needs to verify:

Requirement 1. No returned strings are tampered with, i.e., $\forall s_i \in R, i \in [1, k]$, $s_i \in D$;

Requirement 2. $\forall i \in [1, k-1]$, $DST(s_i, s_q) \leq \theta$, and $DST(s_i, s_q) \leq DST(s_{i+1}, s_q)$;

Requirement 3. No genuine top-k results are missing, i.e., $\forall s_i \notin R$, $DST(s_i, s_q) \geq DST(s_k, s_q)$.

Requirement 1 can be easily verified by restoring the digests of the root signature of the MB-tree. Requirement 2 can be verified by two steps: (1) verifying the soundness of R by either our VS^2 or E-VS^2 approach; (2) re-compute $DST(s_i, s_q)$ for each $s_i \in R$, and sort the strings of R by their distance in an ascending order. Requirement 3 can be verified by checking whether for all the false fits in the VO, their distance to s_q is longer than $DST(s_k, s_q)$. This is equivalent to checking whether the server returns sound the complete results with regard to the similarity threshold $\theta = DST(s_k, s_q)$. Therefore, we can use VS^2 and E-VS^2 approaches to construct VO and do verification, by using the similarity threshold as $DST(s_k, s_q)$.

8. EXPERIMENTS

In this section, we report the experiment results.

8.1 Experiment Setup

Datasets and queries. We use two real-world datasets collected by *US Census Bureau* in 1990[1]: (1) the *LastName* dataset that contains 88799 last names. The maximum length of a name is 13, while the average is 6.83; and (2) the *FemaleName* dataset including 4475 actor names. The maximum length of a name is 11, while the average length is 6.03. We designed ten single-string similarity search queries. For the following results, we report the average of the ten queries.

Parameter setup. The parameters include: (1) string edit distance threshold θ, (2) dimension d of the embedding space, and (3) the fanout f of MB-tree nodes (i.e., the number of entries that each node contains). The details of the parameter settings can be found in Table 2 in Appendix. We also include the details of the query selectivity in Table 3 in Appendix.

[1] `http://www.census.gov/topics/population/genealogy/data/1990_ce`

Due to space limits, the details of experimental environments are in Appendix. For the following discussions, we use the following notations: (1) n: the number of strings in the dataset, (2) n_R: the number of similar strings, (3) n_{MF}: the number of MFs, (4) n_C: the number of C-strings, (5) n_F: the number of FP-strings, (6) n_{DBH}: the number of DBHs, and (7) n_{DS}: the number of DBH-strings. We use σ_S to indicate the average string size, and σ_M (σ_S, resp.) as the size of MB-tree node (DBH, resp.).

8.2 VO Construction Time

We measure the VO construction time at the server side by both VS^2 and E-VS^2 methods.

The impact of θ. In Figure 3 (a), we show the VO construction time with regard to different θ values. First, we observe that there is an insignificant growth of VO construction time by VS^2 when θ increases. For example, on the *LastName* dataset, the time increases from 0.7085 seconds to 0.7745 seconds when θ changes from 2 to 6. This is because the number of MFs reduces with the increase of θ. For example, when $\theta = 2$, $n_{MF} = 390$, while $\theta = 6$, $n_{MF} = 0$. Consequently, VS^2 visits more MB-tree nodes to construct the VO. However, compared with n, the increase of n_{MF} is not significant. Therefore, the VO construction time of VS^2 slightly increases with the growth of θ. Second, we observe the dramatic decrease in VO construction of the E-VS^2 when θ increases. This is because with the increase of θ value, n_R rises, while n_{DS} reduces sharply (e.g., on *LastName* dataset, when $\theta = 2$, $n_R = 136$, $n_{DS} = 13928$; when $\theta = 3$, $n_R = 1104$, $n_{DS} = 3568.6$. Since the complexity of VO construction is cubic to n_{DS}, the total VO construction time decreases intensively when n_{DS} decreases.

The impact of d. We change the dimension d of the embedding space and observe its impact on VO construction time. The results are displayed in Figure 3 (b). On both datasets, d has no effect on VS^2, because VS^2 does not interact with the embedding space. However, the time performance of E-VS^2 increases with the d value. Intuitively, larger dimension leads to smaller n_F, and thus larger n_{DS}. As the VO construction time of E-VS^2 is cubic to n_{DS}, the VO construction time increases when d increases.

The impact of f. For both approaches, we observe that the VO construction time is stable for various f fanout values (Figure 3 (c)). This is because the complexity of VO construction is decided by the number of strings in the dataset and the number of DBH-strings. Both numbers do not change by f values.

8.3 VO Size

We measure the size of the VO constructed by the VS^2 and E-VS^2 approaches.

The impact of θ. The results are shown in Figure 4 (a). The first observation is that the VO size of VS^2 with θ value. Apparently, larger θ values lead to more similar strings (i.e., larger n_R), fewer MFs (i.e., smaller n_{MF}), and fewer C-strings (i.e., smaller n_C). As the VO size is decided by $(n_R+n_C)\sigma_S+n_{MF}\sigma_M$, where $\sigma_M/\sigma_S \approx 10$, and the increase of $(n_R + n_C)\sigma_S$ is cancelled out by the decrease of $n_{MF}\sigma_M$, the VO size of VS^2 approach stays relatively stable. On the contrary, for the E-VS^2 approach, since $\sigma_D >> \sigma_S, \sigma_M$, the VO size is dominately decided by n_{DBH}. We observe that the slight increase of θ values lead to sharp decrease of n_{DS} and thus n_{DBH} (e.g., when $\theta = 2$, $n_{DBH} = 1272$; when $\theta = 6$, $n_{DBH} = 5$). Thus the VO size decreases significantly for larger θ. When $\theta \geq 4$, the VO size of E-VS^2 is very close to that of VS^2.

The impact of d. From the results reported in Figure 4 (b), we observe that VO size of VS^2 is not affected by various d values. This is straightforward as VS^2 does not rely on embedding. On the other hand, the VO size of E-VS^2 increases with larger dimension value, since larger dimension leads to smaller n_F, larger n_{DS}, and thus larger n_{DBH}. Furthermore, when d increases, the average DBH size increases too. These two factors contribute to the growth of VO size for the E-VS^2.

The impact of f. As shown in Figure 4 (c), the VO size decreases with the growth of the fanout f. First, for VS^2, recall that its VO size is calculated as $(n_R+n_C)\sigma_S+n_{MF}\sigma_M$. When f increases, n_R is unchanged. Meanwhile, n_C slightly grows with f (e.g., when $f = 10$, $n_C = 86,216$, while when $f = 40$, $n_C = 87,695$). Furthermore, when f increases, n_{MF} decreases (when $f = 10$, $n_{MF} = 148$, while when $f = 40$, $n_{MF} = 0$). Also $\sigma_M/\sigma_S \approx 10$. Therefore, the decrease of n_{MF} leads to smaller VO size. For E-VS^2, by which the VO size is measured $(n_R + n_F)\sigma_S + n_{MF}\sigma_M + n_{DBH}\sigma_D$, again n_R is unchanged, while n_{MF} decreases for larger f. However, n_F and n_{DBH} keep relatively stable with different f (e.g., when $f = 10$, $n_F = 82688$, $n_{DBH} = 80$; when $f = 40$, $n_F = 84108$, $n_{DBH} = 80$). Therefore, the VO size by E-VS^2 approach decreases with the growth of f.

Another observation is that the VO size of VS^2 is always larger than that of E-VS^2. This is straightforward as E-VS^2 has to include DBHs in VO, which contributes to a substantial portion of VO in terms of its size. Nevertheless, as n_{DBH} is small in most cases (always smaller than 100), the additional VO size required by E-VS^2 is not substantial compared with the total VO size.

8.4 VO Verification Time

In this section, we measure the VO verification time at the client side. We split the verification into five components: (1) time to verify the similarity of the returned similar strings (Result Verification), (2) time to re-construct the root signature of MB-tree and verifying NC-strings (MBtree Verification), (3) time to compute edit distance for C-strings (C-string Verification), (4) time to calculate the Euclidean distance for DBH-strings (DBH-string Verification), and (5) time to compute the edit distance for FP-strings (FP-string Verification). We use the *FemaleName* dataset and report these five components in details.

Before we discuss specific parameters, an important observation is that the verification time of E-VS^2 can be as small as 75% of VS^2. This proves that E-VS^2 can save the verification time at the client side significantly.

The impact of θ. From Figure 5 (a), we observe that, first, the MBtree Verification time keeps stable. This is because the number of MBHs is small (always smaller than 10). Even though the increase of θ decreases the number of MBHs, the time to re-construct the root signature does not increase much. Second, the Result Verification time (0.4% - 32% of the total verification time) increases sharply with θ as the number of similar strings increases fast with θ. For example, when $\theta = 2$, $n_R = 31$, while when $\theta = 6$, $n_R = 2385$. Third, the C-string Verification time decreases when θ grows, since the number of C-strings drops fast with the rapid growth of similar strings. For VS^2, there is no

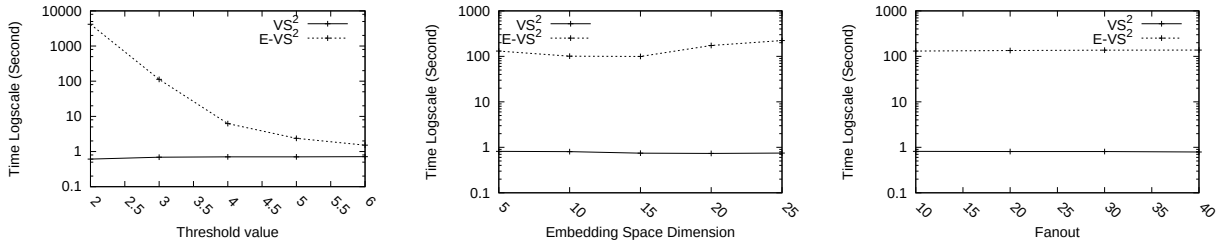

(a) Various θ ($d = 5, f = 10$) (b) Various d ($\theta = 3, f = 10$) (c) Various f ($\theta = 3, d = 5$)

Figure 3: VO construction time (*LastName* dataset)

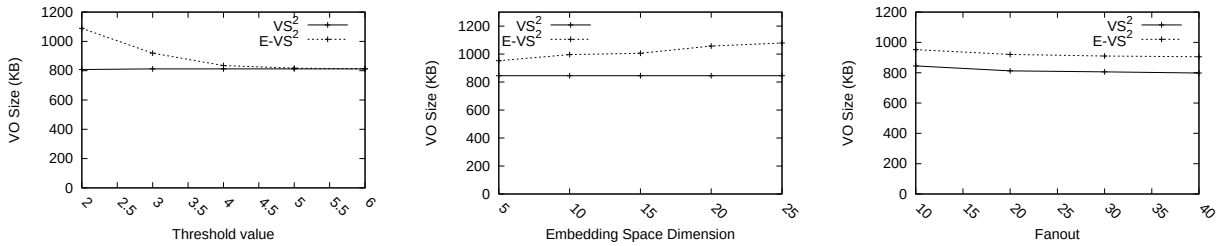

(a) Various θ ($d = 5, f = 10$) (b) Various d ($\theta = 3, f = 10$) (c) Various f ($\theta = 3, d = 5$)

Figure 4: VO size (*LastName* dataset)

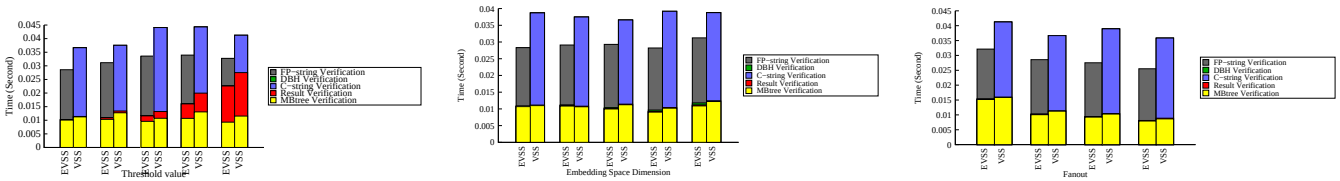

(a) Various θ ($d = 5, f = 10$) (b) Various d ($\theta = 2, f = 10$) (c) Various f ($\theta = 2, d = 5$)

Figure 5: VO verification time (*FemaleName* dataset)

$DBH-$ or FP-string, leading to zero DBH-string and FP-string Verification time. Regarding the E-VS^2 approach, we only discuss DBH-string Verification time and FP-string Verification time, as the other components are the same as the VS^2 approach. Due to the efficient Euclidean distance calculation, the DBH-string Verification time is very small (smaller than 0.2% of the total verification time). The FP-string Verification time decreases when θ grows, as a large portion of FP-strings become similar when θ increases. Overall, the total verification time of $E-VS^2$ increases when θ changes from 2 to 4, but keeps stable after that. This is because when θ increases from 2 to 4, n_{DS} drops very fast (from 1349 to 45). The time saved by verifying the DBHs thus shrinks. When $\theta > 4$, n_{DS} does not change much (from 45 to 4). Thus the total verification time keeps stable.

The impact of d. We only discuss the verification time of E-VS^2 as the verification time of VS^2 does not rely on the dimension of the embedding space. According to the results shown in Figure 5 (b), for E-VS^2, the total verification time keeps stable with the increase of d. The reason is that the number of DBH-strings varies little much with the increase of d. When $d = 5$, $n_{DS} = 1349$; when $d = 25$, $n_{DS} = 1357$. This shows that $E - VS^2$ is efficient even for the high-dimension embedding space.

The impact of f. According to the results shown in Figure 5 (c), larger f value results in shorter verification time for both approaches. The reason is that larger f leads to a smaller MB-tree, and thus small time to re-compute the tree's root signature.

9. CONCLUSION

In this paper, we designed two efficient authentication methods, namely VS^2 and E-VS^2, for outsourced string similarity search. Both VS^2 and E-VS^2 approaches are based on a novel authentication data structure named MB-tree that integrates both B^{ed}-tree and Merkle hash tree. The E-VS^2 approach further applies string embedding methods to merge dissimilar strings into smaller VO. Experimental results show that our methods authenticate similarity query searches efficiently. In the future, we will investigate how to design authentication methods for database with updates. We also plan to study how to authenticate the correctness of *privacy-preserving* string similarity search.

10. APPENDIX

10.1 Proof of Theorem 5.4

We construct a graph $G = (V, E)$ such that for each point $P_i \in \mathcal{P}$, it corresponds to a vertex $v_i \in V$. For any two vertices v_i and v_j that correspond to two points P_i and P_j, there is an edge $(v_i, v_j) \in E$ if $dst_{min}(P_q, R) > \theta$, where R is the MBH of P_i and P_j. We have the following theorem: given the graph $G = (V, E)$ constructed as above, for any clique C in G, let R be the MBH constructed from the points corresponding to the vertice in C. Then R must be a DBH.

Proof. We prove it by induction. Let $|C|$ denote the number

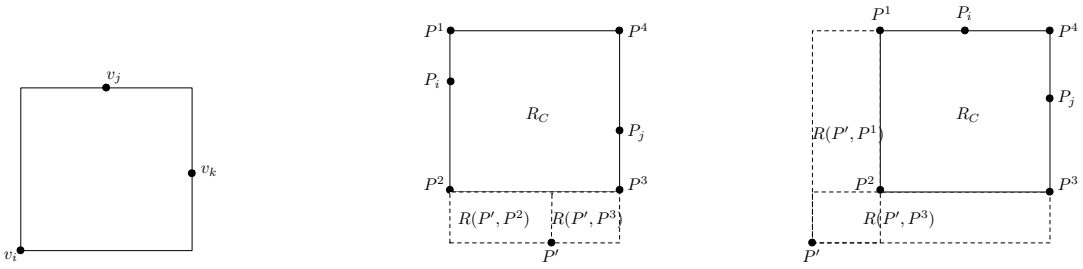

(a) Base case (b) Case 2 of induction (c) Case 3 of induction

Figure 6: Illustration: Proof of Theorem 5.4

of vertices in the clique. It is trivial to show that the theorem holds if $|C| < 3$. Next, we mainly discuss $|C| \geq 3$.

Base case. When $|C| = 3$, let v_i, v_j and v_k denote the vertices in the clique C. Let R_{ij}, R_{jk}, and R_{ik} be the MBHs constructed from the pairs (v_i, v_j), (v_j, v_k), and (v_i, v_k) respectively. Let R_{ijk} be the MBH constructed from v_i, v_j, and v_k. Apparently $R_{ijk} = R_{ij} \cup R_{jk} \cup R_{ik}$. Given the fact that $dst_{min}(p, R_{ij}) > \theta$, $dst_{min}(p, R_{ik}) > \theta$, and $dst_{min}(p, R_{jk}) > \theta$, it must be true that $dst_{min}(p, R_{ijk}) > \theta$. Therefore, R must be a DBH.

Induction step. If we add v' into C, we get a new clique C'. Let $R_{C'}$ be the MBH constructed from C'. Next, we prove that $R_{C'}$ is always a DBH. We prove this for three cases: (1) $P' \in R_C$, (2) P' falls out of the range of R_C at one dimension, and (3) P' falls out of the range of R_C at both dimensions.

Case 1. $P' \in R_C$. This case is trivial as it is easy to see that $R_{C'} = R_C$. So $R_{C'}$ must be a DBH.

Case 2. At exactly one dimension, P' falls out of R_C. Then it must be true that either $P'[i] < l_i^C$ or $P'[i] > u_i^C$, for either $i = 1$ or $i = 2$. Without loss of generality, we define the four boundary nodes of R_C as P^1, P^2, P^3, and P^4 (as shown in Figure 6 (b)). Also we assume that $P'[1] \in [l_1^C, u_1^C]$ and $P'[2] < l_2^C$. It is easy to see that $R_{C'} = (< l_1^C, u_1^C > , < P'[2], u_2^C >) = R_C \cup R(P', P^2) \cup R(P', P^3)$. Apparently, R_C is covered by $R_{C'}$.

Before we prove that $R_{C'}$ is a DBH, we present a lemma.

LEMMA 10.1. *Given two rectangles $R_1(< l_1^1, u_1^1 >, \ldots, < l_d^1, u_d^1 >)$ and $R_2(< l_1^2, u_1^2 >, \ldots, < l_d^2, u_d^2 >)$ in the same Euclidean space, if R_1 is covered by R_2, i.e. $l_i^2 \leq l_i^1 \leq u_i^1 \leq u_i^2$ for any $i = 1, \ldots, d$, then $dst_{min}(P, R_1) \geq dst_{min}(P, R_2)$ for any point P.*

Lemma 10.1 states that if R_1 is covered by R_2, for any point P, its minimum distance to R_1 is no less than the minimum distance to R_2.

Next, let's consider $R_{C'}$ that is constructed from adding the point P' to the existing DBH R_C. We pick a point $P_i \in R_{C'}$ $(1 \leq i \leq t)$ s.t. $P_i[1] = l_1^C$ and $P_i[2] \in [l_1^C, u_1^C]$. Apparently, $R(P', P^2)$ is covered by $R(P', P_i)$. Because there is an edge between v' and v_i in the graph G, it must be true that $dst_{min}(P_q, R(P', P_i)) > \theta$. Following Lemma 10.1, we can infer that $dst_{min}(P_q, R(P', P^2)) > \theta$. Similarly, there must be a point P_j with $P_j[1] = u_1^C$ and $P_j[2] \in [l_1^C, u_1^C]$. Because $R(P', P^3)$ is covered by $R(P', P_j)$ and $dst_{min}(P_q, R(P', P_j)) > \theta$, we can prove that $dst_{min}(P_q, R(P', P^3)) > \theta$. Thus, we prove that $dst_{min}(P_q, R_{C'}) > \theta$ and $R_{C'}$ is a DBH.

Case 3. On both dimensions, P' falls out of R_C. Formally, either $p'[i] < l_i^C$ or $P'[i] > u_i^C$, for both $i = 1, 2$. Without loss of generality, we assume that $P'[0] \in [l_1^C, u_1^C]$ and $P'[2] < l_2^C$. It is easy to see that $R_{C'} = (< P'[1], u_1^C > , < P'[2], u_2^C >) = R_C \cup R(P', P^1) \cup R(P', P^3)$. There must exists a $P_i (1 \leq i \leq t)$ s.t. $P_i[2] = u_2^C$ and $P_i[1] \in [l_1^C, u_1^C]$. In other words, $R(P', P^1)$ is covered by $R(P', P_i)$. As $dst_{min}(P_q, R(P', P_i)) > \theta$, it must be true that $dst_{min}(P_q, R(P', P^1)) > \theta$. Similar to Case 2, we can prove that $dst_{min}(P_q, R(P', P^3)) > \theta$ based on P_j. Thus, we prove that $dst_{min}(P_q, R_{C'}) > \theta$ and $R_{C'}$ is a DBH.

10.2 Special Case of $MDBH$ Problem

There is a special case where the $MDBH$ problem can be solved in polynomial time. In particular, when the embedded points of all DBH-strings lie on a single line, we can construct a minimal number of DBHs in the complexity of $O(\ell)$, where ℓ is the number of DBH-strings. Let $\mathcal{L}$ be the line that the embedded points of DBH-strings lie on. We draw a perpendicular line from P_q to $\mathcal{L}$. Let $dst(P_q, \mathcal{L})$ be the distance between P_q and $\mathcal{L}$. Depending on the relationship between θ and $dst(P_q, \mathcal{L})$, there are two cases:

Case 1: $dst(P_q, \mathcal{L}) > \theta$. We construct the MBH of all the embedded points of DBH-strings.

Case 2: $dst(P_q, \mathcal{L}) \leq \theta$. The perpendicular line splits all points of DBH strings into two subsets, P_L and P_R, where P_L includes the embedded points that are at one side of $\mathcal{L}$, and P_R be the points at the other side. A special case is that P_q lies on $\mathcal{L}$. For this case, P_q still splits all points on $\mathcal{L}$ into two subsets, P_L and P_R. It is possible that P_L or P_R is empty. For each non-empty P_L or P_R, we construct a corresponding MBH.

We have the following theorem.

THEOREM 10.1. **Proof** *For Case 1, if $dst(P_q, \mathcal{L}) > \theta$, for any point P located on $\mathcal{L}$, it must be true that $dst(P_q, P) \geq dst(P_q, \mathcal{L}) > \theta$. For Case 2, if P_L is non-empty, then $dst_{min}(P_q, P_L) \geq$*

$min\{dst(P_q, P)|P \in P_L\} > \theta$. *So* P_L *must be a* DBH. *The same reasoning holds for* P_R.

10.3 Experiments: Experimental environment

We implement both VS^2 and E-VS^2 approaches in C++. The hash function we use is the *SHA256* function from the OpenSSL library. We execute the experiments on a machine with 2.5 GHz CPU and 6 GB RAM, running Mac OS X 10.10.

10.4 Experiments: Parameter Settings

Table 2 includes the details of the parameter settings of our experiments.

10.5 Experiments: Selectivity of Search Queries

In Table 3, we report the selectivity of the threshold values on both datasets, where selectivity is defined as the percentage of similar strings in the dataset. The reported result is the average selectivity of 10 query strings.

θ	*Lastname*	*Femalename*
2	0.1378	0.7277
3	1.1188	3.329
4	5.214	10.61
5	16.4	27.49
6	39.015	55.80

Table 3: Selectivity (%) of search queries w.r.t. different threshold values

11. REFERENCES

[1] A. Arasu, V. Ganti, and R. Kaushik. Efficient exact set-similarity joins. In *Proceedings of the 32nd international conference on Very large data bases (VLDB)*, pages 918–929, 2006.

[2] S. Chaudhuri, K. Ganjam, V. Ganti, and R. Motwani. Robust and efficient fuzzy match for online data cleaning. In *Proceedings of the 2003 ACM SIGMOD international conference on Management of data*, pages 313–324, 2003.

[3] W. Cheng and K.-L. Tan. Query assurance verification for outsourced multi-dimensional databases. *Journal of Computer Security*, 17(1):101–126, 2009.

[4] B. Chor, E. Kushilevitz, O. Goldreich, and M. Sudan. Private information retrieval. *Journal of the ACM (JACM)*, 45(6):965–981, 1998.

[5] D. Comer. Ubiquitous b-tree. *ACM Computing Surveys (CSUR)*, 11(2):121–137, 1979.

[6] G. Cong, C. S. Jensen, and D. Wu. Efficient retrieval of the top-k most relevant spatial web objects. *Proceedings of the VLDB Endowment*, 2(1):337–348, 2009.

[7] S. Eidenbenz and C. Stamm. Maximum clique and minimum clique partition in visibility graphs. In *Theoretical Computer Science: Exploring New Frontiers of Theoretical Informatics*, pages 200–212. Springer, 2000.

[8] C. Faloutsos and K.-I. Lin. *FastMap: A fast algorithm for indexing, data-mining and visualization of traditional and multimedia datasets*, volume 24. ACM, 1995.

[9] R. Gennaro, C. Gentry, and B. Parno. Non-interactive verifiable computing: Outsourcing computation to untrusted workers. In *Advances in Cryptology (CRYPTO)*, pages 465–482. 2010.

[10] M. T. Goodrich, C. Papamanthou, D. Nguyen, R. Tamassia, C. V. Lopes, O. Ohrimenko, and N. Triandopoulos. Efficient verification of web-content searching through authenticated web crawlers. *Proceedings of the VLDB Endowment*, 5(10):920–931, 2012.

[11] H. Hacigümüş, B. Iyer, C. Li, and S. Mehrotra. Executing sql over encrypted data in the database-service-provider model. In *Proceedings of the 2002 ACM SIGMOD international conference on Management of data*, pages 216–227, 2002.

[12] G. Hjaltason and H. Samet. Contractive embedding methods for similarity searching in metric spaces. Technical report, Technical Report TR-4102, Computer Science Department.

[13] G. R. Hjaltason and H. Samet. Properties of embedding methods for similarity searching in metric spaces. *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, 25(5):530–549, 2003.

[14] R. Housley, W. Ford, W. Polk, and D. Solo. Internet X.509 Public Key Infrastructure Certificate and CRL Profile. RFC 2459, RFC Editor, 1999.

[15] L. Hu, W.-S. Ku, S. Bakiras, and C. Shahabi. Spatial query integrity with voronoi neighbors. *IEEE Transactions on Knowledge and Data Engineering*, 25(4):863–876, 2013.

[16] L. Jin, C. Li, and S. Mehrotra. Efficient record linkage in large data sets. In *Database Systems for Advanced Applications, 2003.(DASFAA 2003). Proceedings. Eighth International Conference on*, pages 137–146. IEEE, 2003.

[17] N. Koudas, S. Sarawagi, and D. Srivastava. Record linkage: similarity measures and algorithms. In *Proceedings of the 2006 ACM SIGMOD international conference on Management of data*, pages 802–803, 2006.

[18] C. Li, J. Lu, and Y. Lu. Efficient merging and filtering algorithms for approximate string searches. In *IEEE 24th International Conference on Data Engineering (ICDE)*, pages 257–266, 2008.

[19] F. Li, M. Hadjieleftheriou, G. Kollios, and L. Reyzin. Dynamic authenticated index structures for outsourced databases. In *Proceedings of the 2006 ACM SIGMOD international conference on Management of data*, pages 121–132, 2006.

[20] F. Li, M. Hadjieleftheriou, G. Kollios, and L. Reyzin. Dynamic authenticated index structures for outsourced databases. In *Proceedings of the 2006 ACM SIGMOD international conference on Management of data*, pages 121–132. ACM, 2006.

[21] F. Li, M. Hadjieleftheriou, G. Kollios, and L. Reyzin. Authenticated index structures for aggregation queries. *ACM Transaction of Information and System Security*, 13:1–35, December 2010.

[22] A. J. Menezes, P. C. Van Oorschot, and S. A. Vanstone. *Handbook of applied cryptography*. CRC press, 1996.

[23] E. Mykletun, M. Narasimha, and G. Tsudik.

Parameter	dataset	setting
Similarity threshold θ	*Lastname*	2, 2.5, 3. 3.5, 4, 4.5, 5, 5.5, 6
	Femalename	2, 3, 4, 5, 6
Dimension d of embedding space	*Lastname*	5, 10, 15, 20, 25
	Femalename	5, 10, 15, 20, 25
MB-tree fanout f	*Lastname*	10, 15, 20, 25, 30, 35, 40
	Femalename	5, 10, 15, 20

Table 2: Parameter settings

Authentication and integrity in outsourced databases. *Transaction of Storage*, 2:107–138, May 2006.

[24] H. Pang and K. Mouratidis. Authenticating the query results of text search engines. *Proceedings of the VLDB Endowment*, 1(1):126–137, 2008.

[25] H. Pang, J. Zhang, and K. Mouratidis. Scalable verification for outsourced dynamic databases. *Proceedings of VLDB Endowment*, 2:802–813, August 2009.

[26] S. Papadopoulos, L. Wang, Y. Yang, D. Papadias, and P. Karras. Authenticated multistep nearest neighbor search. *Knowledge and Data Engineering, IEEE Transactions on*, 23(5):641–654, 2011.

[27] C. Papamanthou, R. Tamassia, and N. Triandopoulos. Optimal verification of operations on dynamic sets. In *Advances in Cryptology (CRYPTO)*, pages 91–110. 2011.

[28] R.C.Merkle. Protocols for public key cryptosystems. In *Symposium on Security and Privacy*, pages 122–134, 1980.

[29] D. Wu, B. Choi, J. Xu, and C. S. Jensen. Authentication of moving top-k spatial keyword queries. *IEEE Transactions on Knowledge and Data Engineering*, 27(4):922–935, 2015.

[30] Y. Yang, D. Papadias, S. Papadopoulos, and P. Kalnis. Authenticated join processing in outsourced databases. In *Proceedings of the 2009 ACM SIGMOD International Conference on Management of data*, pages 5–18. ACM, 2009.

[31] Y. Yang, S. Papadopoulos, D. Papadias, and G. Kollios. Spatial outsourcing for location-based services. In *Proceedings of IEEE 24th International Conference on Data Engineering (ICDE)*, pages 1082–1091, 2008.

[32] M. L. Yiu, E. Lo, and D. Yung. Authentication of moving knn queries. In *IEEE 27th International Conference on Data Engineering (ICDE)*, pages 565–576, 2011.

[33] Z. Zhang, M. Hadjieleftheriou, B. C. Ooi, and D. Srivastava. Bed-tree: an all-purpose index structure for string similarity search based on edit distance. In *Proceedings of the 2010 ACM SIGMOD International Conference on Management of data*, pages 915–926. ACM, 2010.